Product Brief

3rd Gen Intel® Xeon® Scalable Processors in HPC

Up to
1.53x higher
geomean HPC performance
vs prior gen[1]

Up to
1.47x higher
Stream Triad performance
vs prior gen[2]

Up to
1.53x higher
FSI Kernel performance
vs prior gen[3]

Flexible, scalable HPC powers progress

Increasingly, high performance computing (HPC) accelerates innovation in verticals like manufacturing, finance, life sciences, oil and gas, energy, and more. Powering the demands of these diverse HPC workloads requires flexibility, customization, and the ability to scale.

The latest 3rd Gen Intel® Xeon® Scalable processors deliver outstanding performance for challenging real-world HPC applications with significant improvements over the prior gen. Innovations in core architecture, increased memory bandwidth, built-in AI and HPC acceleration, and advanced security technologies enable users to minimize complexity and unleash the value of data. Designed through decades of innovation and optimized for the most demanding workload requirements, on premise and in the cloud, these powerful, future-ready processors deliver:

- **Flexible performance on real-world HPC workloads**
- **Enhanced performance** using Intel's newest processor architecture and 10nm process technology to boost a wide range of HPC applications[4]
- **Increased memory bandwidth,** from six to eight channels, running at faster speeds
- **Flexible configuration** to set cores and frequencies to best match workload requirements with **Intel® Speed Select Technology**

This brief covers two-socket configurations. For more information about 3rd Gen Intel Xeon Scalable processors in four- and eight-socket configurations, visit intel.com/xeonscalable.

Redefining what's possible

HPC plays a vital part in driving critical advancements across a wide range of vertical markets. Several drivers are fueling growth in HPC[5]:

The rise of cloud computing has made HPC more accessible to more users, providing the capability to process large volumes of data, serve more workloads, and run custom optimizations for greater efficiencies.

HPC continues to be a key tool for meeting real-time business challenges and solving the latest science questions. Traditionally it has been used in modeling and simulation, allowing comparison of complex computational models versus observed data, and to speed research and development for products and services. Now, new data-driven paradigms are heavily utilizing HPC to derive deeper insights from ever-increasing amounts of data.

Additionally, the ongoing introductions of new hardware, algorithms, applications and use cases in machine and deep learning are generating interest in HPC systems capable of delivering fast performance for aggressive training, inferencing, and decision support tasks.

Finally, on a global level, more and more, national and regional governments view exascale computing as not only critical for scientific, engineering, and security applications but also as a strategic national resource with important economic and industrial impact on their domestic economies.[6]

The insatiable demand for performance coupled with a need for flexibility today and in the future is raising the bar for HPC servers. There is a growing need for an open architecture that can support and be optimized for evolving and diverse applications—a system with CPUs that can be relied upon to quickly deploy a wide variety of applications while protecting data at-rest, in-flight, and in-use, without compromising privacy.

Flexible performance for real-world applications

At Intel, we know that meeting today's needs—efficiently running different applications with different architectures at various points throughout the day—requires a flexible infrastructure. That's why we've created a platform that is optimized for multi-cloud and AI and capable of serving many workloads, anywhere, at any time.

Available with up to 40 cores per processor in a two-socket configuration, the latest 3rd Gen Intel Xeon Scalable processors deliver up to 46 percent higher geomean general performance than prior-generation processors.[7] New benefits include PCI Express Gen 4 support, increased memory bandwidth to support memory-intensive workloads, memory capacity per processor up to 6TB per processor/socket, and advanced security technologies.

More cores, increased performance

Up to **40** cores per CPU

Up to **1.46x** higher geomean General performance vs prior gen[8]

Intel's **Newest** processor architecture **10nm process** technology

More bandwidth for memory-intensive workloads

8 DDR4 memory channels

vs. 6 in the prior gen

Up to **3200** megatransfers per second

vs. 2933 MT/s in the prior gen

Up to **1.60x** memory bandwidth

in a 2-socket configuration (8 channels, 256GB DDR4) vs. the prior gen (6 channels, 128GB DDR4)

Faster I/O for real-time insights

PCIe Gen 4

Increased I/O capacity

64 lanes PCIe Gen 4 per socket

compared to 48 lanes of PCIe 3.0 in the prior gen

For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.

Scalable, flexible, customizable

Designed through decades of innovation, 3rd Gen Intel Xeon processors provide the flexibility, capability, and interoperability to power diverse workload demands. The platform includes these enhanced Intel technologies:

Intel® Speed Select Technology (Intel® SST): Now with four modes, allows you to tailor the configuration of cores and frequencies to best match diverse workload requirements without requiring a reboot.

Learn more at intel.com/speedselect.

New Q SKU, optimized for liquid-cooled systems: Designed to deliver top-bin performance with high core count and frequencies.

Intel® Optane™ Persistent Memory 200 Series: Offers capacities up to 512GB per module using Intel Optane Persistent Memory 200 Series and up to 6TB total system memory.

Learn more at intel.com/optanepersistentmemory.

Intel® Optane™ SSD: The world's fastest data center SSD[9], the Intel Optane SSD P5800X delivers exceptional HPC performance with predictable low latencies, unprecedented endurance, and PCIe Gen 4 support—a powerful solution for hot-tier caching and accelerating all types of storage solutions.

Learn more at intel.com/optane.

In addition, connectivity, storage, memory, FPGA solutions, and cross-architecture software tools delivered by Intel can enable HPC users to move more, store more, and process everything. Together, these technologies and 3rd Gen Intel Xeon Scalable processors provide an effective and efficient platform performance for outstanding utility, predictability, and peace of mind.

Built-in AI and HPC acceleration for demanding workloads

The balanced architecture includes optimizations such as built-in acceleration, allowing organizations to deploy most of their AI/ML (artificial intelligence/machine learning) workloads on the Intel Xeon platforms they know and trust—not on expensive, proprietary add-in hardware.

Built-in AI and HPC acceleration

Up to **1.56X** higher AI Inference performance with Intel® Deep Learning Boost vs prior gen[10]

Up to **1.60X** higher geomean Life and Material Science performance vs prior gen[11]

Advanced security technologies

3rd Gen Intel Xeon Scalable processors for HPC deliver a revolutionary step forward for privacy and security.

An added layer of protection

Help protect data and application code in real time, from the edge to the data center and multi-tenant public cloud. **Intel® Software Guard Extensions (Intel® SGX)** enable enhanced collaboration using shared data—without compromising privacy. Intel® SGX is the most tested, researched, and deployed Trusted Execution Environment (TEE) in the data center with the smallest attack surface within the system.

Learn more at intel.com/sgx.

Boost the performance of encryption-intensive workloads, including SSL web serving, 5G infrastructure, and VPN/firewalls with **Intel® Crypto Acceleration**, which enables mitigation of the performance impact of pervasive encryption.

Enhance resistance to physical attacks on data with full memory encryption support from **Intel® Total Memory Encryption**.

A trusted foundation

Help protect, detect, correct, providing NIST SP800-193 firmware resiliency, with Intel® Platform Firmware Resilience (Intel® PFR) using an Intel® FPGA. Platform firmware can be validated before execution, while runtime monitoring and filtering help protect against manipulation. In case of a compromise, Intel PFR provides automated recovery in minutes.

Decades of code and deployment support

Intel is committed to software investments across a spectrum of areas—from libraries to AI frameworks, compilers, parallel programming, open source, and more. These initiatives yield tools that help the ecosystem and developers ensure applications take advantage of hardware capabilities for easy deployments to new platforms. Resources include cross-architecture toolkits to optimize code and accelerate application performance across Intel XPU architectures (CPU, GPU, FPGA).

Intel® oneAPI Toolkits

Intel® oneAPI Toolkits are built on Intel's rich heritage of developer tools to accelerate compute on Intel multi-architectures. To help accelerate workloads and bring out the full value of the new 3rd Gen Intel Xeon Scalable platform's performance, AI and encryption capabilities, software developers can optimize their applications using Intel's oneAPI cross-architecture toolkits.

- **Deliver High-Performance Compute and Cryptography**—The compilers and libraries in the Intel® oneAPI Base and HPC Toolkits automatically take advantage of Intel® AVX-512 and Intel® Deep Learning Boost technology to deliver high-performance applications.
- **Accelerate AI**—With the Intel® oneAPI AI Analytics Toolkit and Intel® Distribution of OpenVINO™ toolkit working together with Intel-optimized open source frameworks, boost AI performance for training and deep learning inference across Intel® architectures from edge to cloud.

Learn more at intel.com/oneAPI

Intel® oneAPI Base Toolkit
Native Code Developers
A core set of high-performance tools for building C++, Data Parallel C++, oneAPI library-based, & Python applications

Specialized Workloads

High-Performance HPC Applications

End-to-end Machine Learning & Deep Learning

Advanced Ray Tracing & Rendering

Data Scientists & AI Developers

OpenVINO

High-performance Inference & Vision

Streamlined deployment, faster time to value

Trusted and verified Intel® Select Solutions for HPC offer an easy and quick-to-deploy infrastructure optimized to help accelerate time to breakthrough, actionable insights, and new product designs.

Learn more about Intel Select Solutions for High Performance Compute.

intel SELECT SOLUTIONS

Driving innovation through long-standing industry collaborations

Through close partnerships with the world's software leaders and solution providers, 3rd Gen Intel Xeon Scalable processors have been optimized for many workload types and performance levels, all with the consistent, open Intel architecture you know and trust.

3rd Gen Intel® Xeon® Scalable processors speed business insights and scientific discoveries

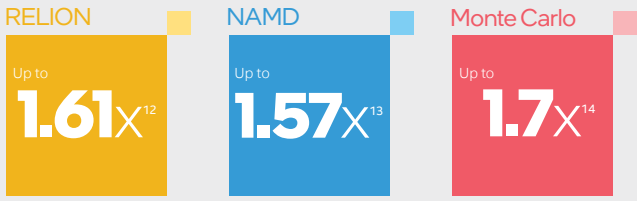

Higher RELION, NAMD, and Monte Carlo performance vs prior gen.

Discover new ways to extract better performance from a wide range of workloads and ensure HPC applications perform at their best. Learn more at intel.com/hpc.

intel XEON PLATINUM

intel XEON GOLD

For more information about	Visit
PROCESSORS	
Intel® Xeon® Scalable Processors	intel.com/xeonscalable
Product Specifications	ark.intel.com
PROCESSOR-RELATED TECHNOLOGIES	
Intel® Advanced Vector Extensions 512 (Intel® AVX-512)	intel.com/avx512
Intel® Speed Select Technology (Intel® SST)	intel.com/speedselect
Intel® Software Guard Extensions (Intel® SGX)	intel.com/sgx
Intel® Resource Director Technology (Intel® RDT)	intel.com/intelrdt
Intel® Artificial Intelligence	ai.intel.com
MEMORY, STORAGE AND CONNECTIVITY	
Intel® Optane™ Persistent Memory	intel.com/optanepersistentmemory
Intel® Optane™ SSD	intel.com/optane
Intel® Ethernet Products	intel.com/networking
SOFTWARE DEVELOPMENT TOOLS	
Intel® oneAPI	intel.com/oneAPI
FAST AND EASY INFRASTRUCTURE DEPLOYMENT	
Intel® Select Solutions for High Performance Compute	Intel Select Solutions for High Performance Compute

Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

[1] Up to 1.53x higher geomean HPC performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[2] Up to 1.47x higher Stream Triad performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[3] Up to 1.53x higher FSI Kernel performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[4] Comparing 3rd Gen Intel Xeon Scalable processors to 2nd Gen Intel Xeon Scalable processors.

[5] Hyperion Research, Updated Worldwide HPC Server Market Status and Forecast, 2019-2014, Including Covid-19 Impacts, #HR1.0040.07.08.2020, https://hyperionresearch.com/olp_document/updated-worldwide-hpc-server-market-status-and-forecast-2019-2024-including-covid-19-impacts/.

[6] Hyperion Research, Updated Worldwide HPC Server Market Status and Forecast, 2019-2014, Including Covid-19 Impacts, #HR1.0040.07.08.2020, https://hyperionresearch.com/olp_document/updated-worldwide-hpc-server-market-status-and-forecast-2019-2024-including-covid-19-impacts/.

[7] Up to 1.46x higher geomean general performance with Platinum 8380 vs prior gen. See [125] at www.intel.com/3gen-xeon-config. Results may vary.

[8] Up to 1.46x higher geomean general performance with Platinum 8380 vs prior gen. See [125] at www.intel.com/3gen-xeon-config. Results may vary.

[9] As compared to generally available PCIe Gen4 x4 Enterprise and Data Center industry SSDs. Results may vary.

[10] Up to 1.56x higher AI Inference performance with Platinum 8380 vs prior gen. See [121] at www.intel.com/3gen-xeon-config. Results may vary.

[11] Up to 1.60x higher geomean Life and Material Science performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[12] Up to 1.61x higher RELION performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[13] Up to 1.57x higher NAMD performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

[14] Up to 1.70x higher Monte Carlo performance on Platinum 8380 vs prior gen. See [108] at www.intel.com/3gen-xeon-config. Results may vary.

0321/JW/MIM/PDF

♻ Please Recycle

346546-001